Nonconvex regularization for inverse problems

Rick Chartrand

Los Alamos National Laboratory

September 28, 2006

Outline

Image Reconstruction
- Regularization
- Examples
- Algorithm

The magic of compressed sensing
- Sparse tomography
- Compressed sensing
- Error correction
- 3-D tomography

Conclusions

Image Reconstruction

Image reconstruction can take many forms:

- denoising
- deblurring
- inpainting
- Abel inversion

Each of these is an ill-posed inverse problem.

Regularization

We approach these problems variationally, and deal with the ill-posedness with regularization.

Given image data $\boldsymbol{f}$, find reconstruction $\boldsymbol{u}$ as minimizer of:

$$\int \text{(penalty term)} \qquad + \text{(parameter)} \qquad \int \text{(data-fidelity term)}$$

$$\int \boldsymbol{R(u)} \qquad + \boldsymbol{\lambda} \qquad \int \boldsymbol{DF(P(u), f)}.$$

Penalty term examples

Gaussian smoothing:

$$\int |\nabla u|^2 + \lambda \int |u - f|^2$$

(blurs object edges)

Total-variation regularization:

$$\int |\nabla u| + \lambda \int |u - f|^2$$

(preserves edges, but shortens them)

Nonconvex regularization:

$$\int |\nabla u|^p + \lambda \int |u - f|^2, \quad 0 < p < 1$$

(preserves most object geometries)

Examples

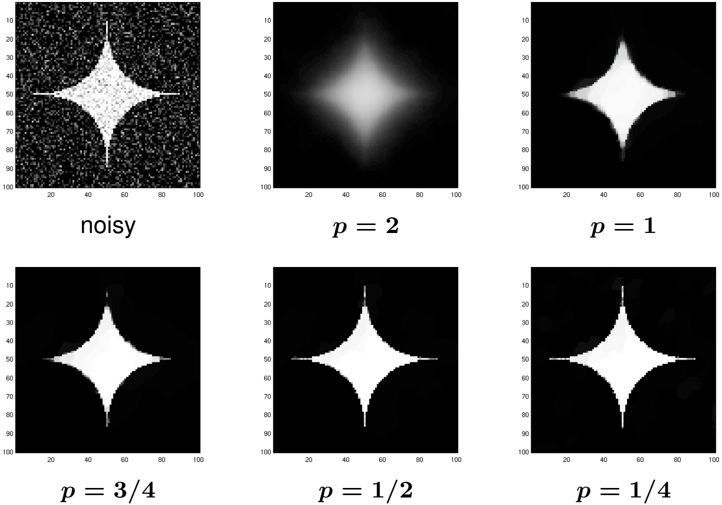

Examples

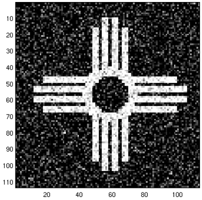

noisy

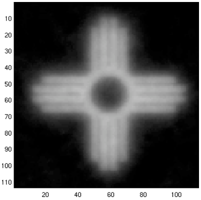

$p = 2$

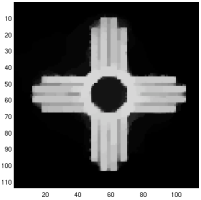

$p = 1$

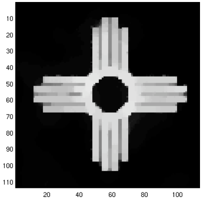

$p = 3/4$

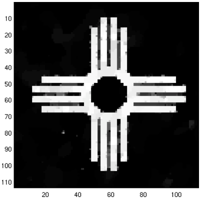

$p = 1/2$

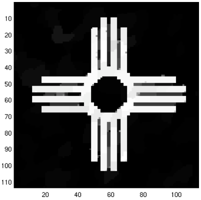

$p = 1/4$

Fixed-point algorithm

Euler-Lagrange equation:

$$\mathbf{0} = -\nabla \cdot \left(|\nabla u|^{p-2} \nabla u\right) + \lambda(u - f).$$

"Lag" the nonlinear portion to get linear system:

$$\mathbf{0} = -\nabla \cdot \left(|\nabla u_n|^{p-2} \nabla u_{n+1}\right) + \lambda(u_{n+1} - f).$$

Converges fast!

Sparse tomography

Suppose we want to reconstruct an image from samples of its Fourier transform. How many samples do we need?

Suppose we have less than 4% of the Fourier transform. Is that enough?

Shepp-Logan phantom

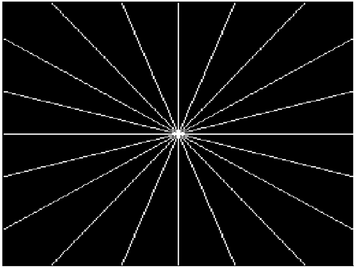

Ω

Nonconvexity again

Yes, using nonconvex minimization:

$$\min_{u} \|\nabla u\|_p, \text{ subject to } \hat{u}|_\Omega = \hat{f}|_\Omega.$$

With $p = 1$, solution is $u = f$ given 17 projections ($\frac{|\Omega|}{|f|} = 6.5\%$).

With $p = 1/2$, 10 projections suffice ($\frac{|\Omega|}{|f|} = 3.8\%$).

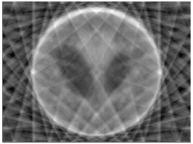

backprojection, 17 views

$p = 1$, 17 views

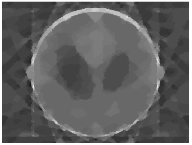

$p = 1$, 10 views

$p = \frac{1}{2}$, 10 views

Compressed sensing

Usual approach to data acquisition and compression:

- acquire the data (all of it)
- compute a sparse representation
- throw away the original data

Problems:

- data may be difficult or expensive to acquire
- dataset may too large to deal with easily

An obvious better way would be to directly acquire a sparse representation, compressed sensing.

Random projections

A general approach is to measure random projections: if the rows of a matrix Φ are M i.i.d. Gaussian vectors, then the solution to

$$\min_{u} \|u\|_p, \text{ subject to } \Phi u = \Phi f,$$

is, for $p = 1$:

- exactly f (with overwhelming probability) if f is K-sparse and $M \geq CK \log N$;
- nearly f if f is nearly K-sparse (i.e., K-compressible) and $M \geq CK \log N$, even if the measurements Φf are noisy

For $p < 1$, we find that fewer measurements are needed to produce the same results.

The geometry of ℓ^p

Why CS works:

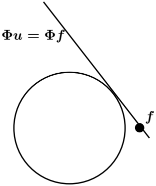

$p = 2$

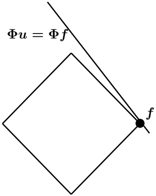

$p = 1$

Why $p < 1$ is better:

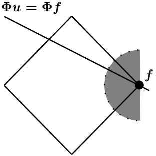

$p = 1$

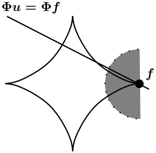

$p < 1$

Los Alamos NATIONAL LABORATORY EST. 1943

Error correction

Let $\boldsymbol{A}$ be a random Gaussian matrix. Can we recover the "plaintext" $\boldsymbol{f}$ if the "ciphertext" $\boldsymbol{Af}$ is corrupted by many, large errors?

If $\boldsymbol{y} = \boldsymbol{Af} + \boldsymbol{e}$, we minimize $\|\boldsymbol{y} - \boldsymbol{A}\tilde{\boldsymbol{f}}\|_{\boldsymbol{p}}$. If $\boldsymbol{e}$ is sparse enough, then the minimizer $\boldsymbol{f}^*$ will be exactly $\boldsymbol{f}$.

Given random $\boldsymbol{B}$ whose kernel is the range of $\boldsymbol{A}$, the problem is equivalent to minimizing $\|\boldsymbol{u}\|_{\boldsymbol{p}}$, subject to $\boldsymbol{Bu} = \boldsymbol{Be}$.

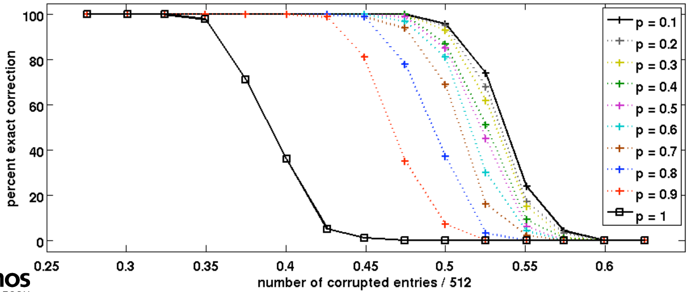

3-D tomography

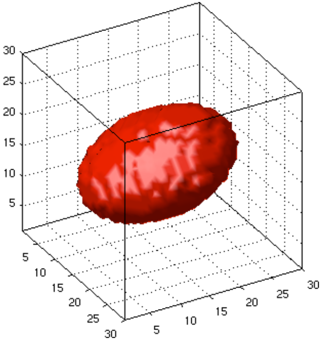

Non-oriented ellipsoid

Six radiographic views suffice for exact reconstruction with $p = 1$, three with $p = 1/2$.

3-D Tomography

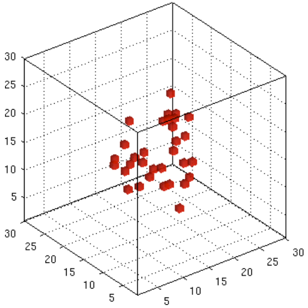

Remove 1% of the voxels, randomly. Four views allow an exact reconstruction of the depleted ellipsoid, to identify defects precisely.

For objects with piecewise-constant density, far less data is needed than for traditional CT methods.

Conclusions

- Nonconvex regularization in image reconstruction improves geometry preservation.
- We have a fast algorithm to do this.
- Compressed sensing is a powerful way to obtain sparse representations from limited data, even more limited in the nonconvex case.
- Current algorithms for nonconvex CS are feasible but not fast.
- The best applications are yet to come.